arXiv:1605.08969v1 [cs.MM] 29 May 2016

Improving Crowdsourced Live Streaming with Aggregated Edge Networks

Chenglei Wu, Zhi Wang, Jiangchuan Liu, Shiqiang Yang
{wcl15@mails., wangzhi@sz., yangshq@}tsinghua.edu.cn, jcliu@cs.sfu.ca

***Abstract*—Recent years have witnessed a dramatic increase of user-generated video services. In such user-generated video services, *crowdsourced* live streaming (e.g., Periscope, Twitch) has significantly challenged today's edge network infrastructure: today's edge networks (e.g., 4G, Wi-Fi) have limited uplink capacity support, making high-bitrate live streaming over such links fundamentally impossible. In this paper, we propose to let broadcasters (i.e., users who generate the video) upload crowdsourced video streams using *aggregated* network resources from multiple edge networks. There are several challenges in the proposal: First, how to design a framework that aggregates bandwidth from multiple edge networks? Second, how to make this framework transparent to today's crowdsourced live streaming services? Third, how to maximize the streaming quality for the whole system? We design a multi-objective and deployable bandwidth aggregation system *BASS* to address these challenges: (1) We propose an aggregation framework transparent to today's crowdsourced live streaming services, using an edge proxy box and aggregation cloud paradigm; (2) We dynamically allocate geo-distributed cloud aggregation servers to enable MPTCP (i.e., multi-path TCP), according to location and network characteristics of both broadcasters and the original streaming servers; (3) We maximize the overall performance gain for the whole system, by matching streams with the best aggregation paths.**

I. Introduction

Recent years have witnessed a dramatic increase of user-generated video services, enabled by the popularity of online video service and online social network service, and their cross pollination [1]. Among such user-generated video services, *crowdsourced* live streaming(e.g., Periscope, Twitch), which allows individuals to broadcast *live* video streams to millions of users on the move, basks in surge of interest. For instance, users collectively watch over 40 years worth of live streaming content on Periscope [2], one of the representative crowdsourced live streaming apps.

In live streaming service, we distinguish *broadcasters* (i.e., users who generate and upload the video streaming) and *viewers* (i.e., users who watch the live streaming). Fig. 1 illustrates the conventional approach for crowdsourced live streaming: the original stream from a broadcaster will be transferred to streaming servers, usually via TCP; then the video stream will be delivered to the viewers by the streaming servers. Although many researchers have been working on this topic, most of these studies [3]–[5] only focused on the video delivery part (i.e., how video data is delivered to users), and left the uploading (i.e., how video data is effectively uploaded to the servers) side unexplored.

Given the fact that today's mobile devices can take high-resolution videos, e.g., iPhone 6s supports 30 fps 4K High Definition video recording[1], it is challenging for today' edge network infrastructure to fulfill the large bandwidth requirement: to effectively online stream a 1080P HD video, the bandwidth required is at least 5 Mbit/s, and 15 Mbit/s for 4K streaming with efficient video compression methods. For live broadcasting which is unable to compress video data efficiently, the bandwidth requirement would be inevitably much larger.

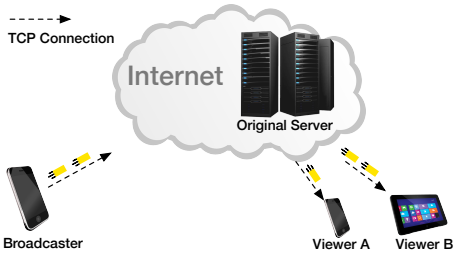

Fig. 1. Crowdsourced live video streaming: users upload live video streaming via edge networks, e.g., 3G.

Wi-Fi networks along with cellular networks (e.g., 3G, 4G) and other networks compose the common edge networks. The bandwidths of current edge networks, however, often fail to satisfy the requirement: almost 60% of users have less than 1 Mbits/s upload bandwidth according to our measurement study on Wi-Fi quality in representative cities. One of the reasons is that local carrier often configures the upload rate much lower than the download rate, since most people have more of a need to download data. And according to Akamai's research, only 21% of U.S. homes have more than 15 Mbits/s bandwidths[2], which also confirms that uplink capacity is relatively limited due to the same reason. The upload capacity of 4G cellular network is insufficient for 4K live streaming and the cost is rather high. Therefore the gap between bandwidth requirement and reality significantly challenged today's wireless infrastructure.

By analyzing a massive data set collected by a crowdsourced Wi-Fi association mobile application which has more than 13 million users, we observed that when users try to connect to a Wi-Fi hotspot, the chance of having more than 10 other Wi-Fi networks around is over 50%. A promising idea to address

[1] According to Apple, a minute of video with 1080p HD at 30 fps is approximately 130MB and a minute of video with 4K higher resolution is approximately 375MB.

[2] http://www.fiercecable.com/story/akamai-only-21-us-homes-have-enough-bandwidth-stream-4k/2015-09-23

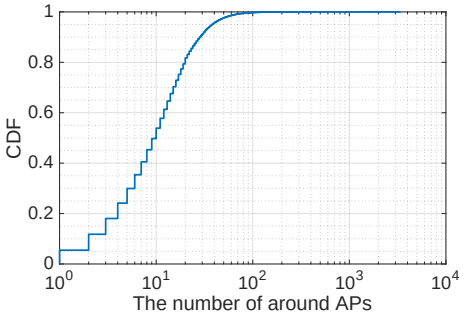

Fig. 2. Number of Wi-Fi access point sensed by the devices.

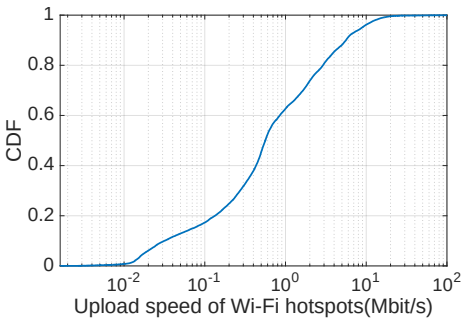

Fig. 3. Upload bandwidth of Wi-Fi hotspots.

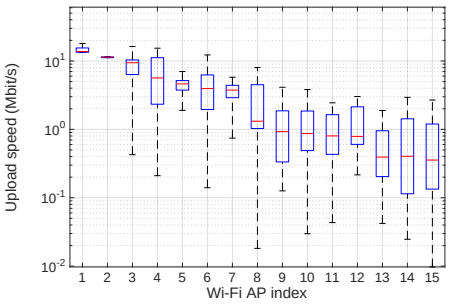

Fig. 4. Upload speeds of Wi-Fi hotspots in the same cellular cell selected from a popular region.

this problem is to utilize all available edge networks simultaneously including Wi-Fi networks and cellular networks. Several previous efforts [6]–[9] have been devoted to using aggregated network resources to improve bandwidth capacity; however, they fail to increase the throughput for applications that originally use single-TCP connections to deliver data, e.g., most of the crowdsourced live streaming services.

In this paper, we propose *BASS* (Bandwidth Aggregation SyStem), a multi-objective and deployable bandwidth aggregation system. The BASS system streams crowdsourced video data using *aggregated* edge network resource via a portable device equipped with multiple network interfaces. BASS dynamically allocates *aggregation servers* that receive video data from users and pass it to the original server, to enable MPTCP (i.e., Multi-path TCP [10]–[13]) for today's representative crowdsourced live streaming platforms.

In our design of the BASS system, we try to fulfill the following requirements. (1) The system should be easy to deploy without any modifications to existing network infrastructure, servers and applications. (2) The design should be able to exploit multiple available edge networks including both Wi-Fi and cellular networks. (3) The strategies should dynamically allocate aggregation servers to provide satisfactory QoS for broadcasters.

Our contributions below provide a system design and strategies that satisfy these requirements.

First, we carry out measurement studies on 13 million users connecting to 17 million Wi-Fi hotspots in a time span of one month, to investigate (1) the availability of Wi-Fi resources in representative metropolitan areas (i.e., more than 50% of the Wi-Fi sessions record that there are over 10 other Wi-Fi hotspots around available), and (2) the insufficiency of most Wi-Fi hotspots' uplink bandwidths (i.e., over 60% of the Wi-Fi sessions have an upload capacity lower than 1 Mbit/s).

Second, we propose a practical aggregation network framework, which makes use of geo-distributed cloud resources for bandwidth aggregation for servers/users located at different places. Today's mobile devices usually do not support using multiple wireless interfaces simultaneously. As opposed to previous study [14] that modifies the client devices, the adoption of middle-box and aggregation server in our design makes no modification to current operating system on mobile devices and makes *BASS* more practical to be deployed for real-world devices and applications. Our design features: (1) The framework works transparently, i.e., broadcasters can use the original apps on their smart phones to upload high-bitrate video streams to the original servers. (2) Dynamical aggregation strategies, including i) a geo-distributed cloud aggregation server allocation algorithm that allocates cloud servers best matching the users and original servers (i.e., large bitrates can be achieved in the aggregation), and ii) a heuristic algorithm to maximize the whole system's overall bandwidth gain.

Third, we implement a prototype on EC2/PlanetLab to verify the effectiveness of our design. In particular, 8 Amazon EC2 and 60 Planet-Lab nodes are employed in our design to evaluate its performance. Compared with traditional TCP based uploading, our design can significantly improve the bandwidth by up to 5 times.

The remainder of this paper is organized as follows. Section II presents our measurement results and motivation. Section III discusses the overall architecture of BASS system. Section IV presents the dynamic allocation algorithm. In Section V, we evaluate our algorithm via both simulations and experiments. Finally, Section VI concludes this paper and provides our plan for future work.

II. Motivation: Measurement Studies

A. Availability of Multiple Wi-Fi Networks

To understand the availability of Wi-Fi resources, in this paper we analyzed a large dataset collected by a popular online crowdsourced Wi-Fi association mobile app in China. This app has over 200 million downloads. It allows users to discover nearby crowdsourced free Wi-Fi hotspots. Every succeed connection is considered as a session, in which QoS statistics (e.g., measure of upload bandwidth of the Wi-Fi AP) will be reported to a trace server. In this paper, we use over 460 million collected sessions of 13 million users connecting to 17 million Wi-Fi hotspots in a time span of one month to evaluate the availability of Wi-Fi resource.

In the traces collected, users report the available Wi-Fi hotspots around, based on which we are able to estimate the number of Wi-Fi hotspots available to a user. Fig. 2 presents the CDF of the number of Wi-Fi hotspots available to users.

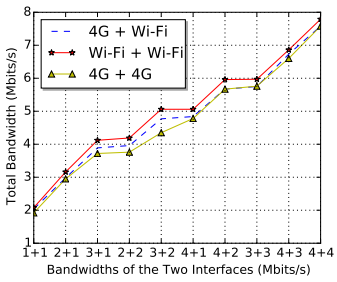

Fig. 5. Effectiveness of MPTCP for upload bandwidth aggregation.

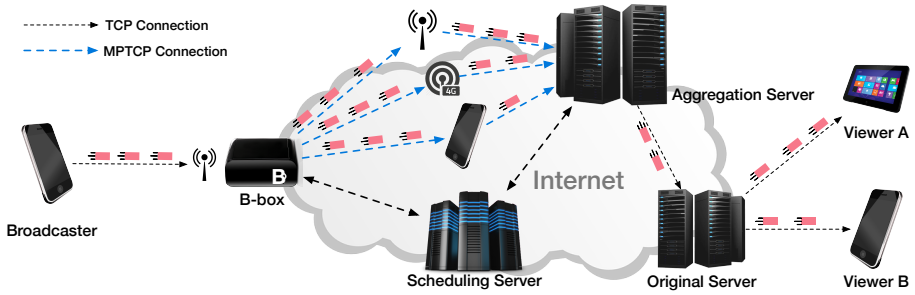

Fig. 6. System architecture of BASS: Videos are uploaded using aggregated edge networks with portable B-box and cloud-based aggregation servers.

We observe that over 50% of these sessions have more than 10 Wi-Fi hotspots around users when they try to connect to a Wi-Fi network. This observation suggests that today's Wi-Fi deployment is highly promising for aggregation usage, i.e., using multiple Wi-Fi networks to aggregation the bandwidth capacity.

In Fig. 3, we present the CDF plot of collected upload bandwidth test results. From this plot, we can observe that almost 60% of collected bandwidth records are lower than 1 Mbit/s, which confirms that the bandwidth of today's edge network are very limited.

We next study the upload bandwidth of Wi-Fi hotspots inside a 100×100m grid. Each sample shows the upload speed distribution of different APs located in the same cellular cell. From Fig. 4, We observe significant difference of the median upload speeds among these APs, varying from 0.4 Mbit/s to 11 Mbit/s. The quality diversity in nearby hotspots suggests scheduling video chunk packets according to the status of the Wi-Fi interfaces is in demand.

B. Bandwidth Improvement via MPTCP

We verify the effectiveness of MPTCP in utilizing multiple wireless interfaces. MPTCP is a modification of the classical TCP that allows end-to-end data traffic to be split across multiple paths, while maintaining TCP connections at the end points (applications). Upper layers of the protocol stack only need to deal with a logical "master" TCP socket. Fig. 5 illustrates a simple experiment conducted to measure the performance of MPTCP. We measured the bandwidth between local client and remote server under different network interface and bandwidth combinations. Both client and server are installed with MPTCP. The network interfaces we used in these experiments are two Wi-Fi interfaces and two 4G cellular interfaces. We observe that MPTCP is able to aggregate bandwidth resources from multiple networks.

To summarize, apart from the constant availability of cellular networks, users often have access to several Wi-Fi networks in urban areas, though these Wi-Fi networks often have limited bandwidth and heterogeneous upload capacities; also, MPTCP aggregation is promising to make full utilization of multiple wireless network resources for crowdsourced live stream uploading.

III. System Components and Architecture of *BASS*

In this section, we provide a detailed description of the BASS system.

A. System Components

The BASS system has three main components. Firstly, a portable device B-box which can be treated as a special mobile Wi-Fi hotspot with several network interfaces and a built-in lightweight linux operating system (e.g., OpenWrt). Secondly, a group of aggregation servers with MPTCP enabled, usually allocated as cloud instances (e.g., Amazon EC2). Finally, a scheduling server which monitors aggregation servers' status and allocates aggregation servers to enable MPTCP.

1) B-box: Bandwidth Aggregation Box: B-box has a built-in lightweight linux operating system, equipped with several types of network interfaces including Wi-Fi and cellular. Thus, B-boxes are able to connect to multiple networks at the same time. The number of each kind of network interface is not fixed and can be adjusted according to user requirement when developed in the real world. B-box is deployed with MPTCP, therefore it can connect to a MPTCP-enabled server via multiple network interfaces. B-box also sets up a Wi-Fi hotspot using one of its Wi-Fi interfaces for users to connect. A B-box runs a proxy program that will receive video chunks uploaded by users via TCP and then upload them to the aggregation server via MPTCP. These proxy programs can be regular proxy programs like Shadowsocks, Squid or protocol converters [15] designed specifically for MPTCP.

2) Aggregation Server: Bandwidth Aggregation in the Cloud: To avoid any modification to the crowdsource live streaming system and make BASS practical for today's crowdsource live streaming systems, we propose to allocate aggregation servers as a middleware to transfer data between original servers and B-boxes via MPTCP. Thus the data between B-box and the aggregation server can be delivered by several sub-flows via multiple paths. With bandwidth aggregation, B-box can obtain an aggregated bandwidth to provide a higher upload capacity for the users.

3) Scheduling Server: Centralized Aggregation Scheduler: Scheduling server is the coordinator between B-boxes and aggregation servers. Aggregation servers report their current

loads (e.g., bandwidth consumed by served broadcasters) to scheduling server regularly. The scheduling server keeps track of the status and resource usages of all aggregation servers. When user starts a B-box, the B-box retrieves available aggregation server list from scheduling server and requests one aggregation server from scheduling server. Scheduling server assigns aggregation server to B-box and decides how much resource on this aggregation server will be allocated.

B. Architecture

Fig. 6 illustrates the architecture of BASS. In a nutshell, a broadcaster starts B-box which automatically requests aggregation server from scheduling server and sets up connection. Broadcaster uploads video streams to B-box. B-box then transmits the data via multiple interfaces to the aggregation server which will continually pass the data to the original server. The original servers are often dedicated streaming servers, for example, RTMP (Real Time Messaging Protocol over HTTP Tunnel) streaming servers deployed by Twitch.tv [4].

Since the original servers and users of a crowdsource live streaming system can be distributed in different locations, we propose to use a geo-distributed cloud solution to allocate aggregation servers at different locations accordingly, such that the aggregated bandwidth and latency can both be improved [16]. To optimize network performance and provide good latency and throughput performance, a deployment limited to 11 locations for North America or a total between 36 and 72 cloud-service locations with good peering connections on a global scale will be sufficient [17].

In our detailed design later, we will present how aggregation servers are allocated and how B-boxes are matched with different aggregation servers.

IV. Detailed Design of *BASS*

In this section, we present how aggregation servers are allocated and how they are matched with different B-boxes.

A. Dynamical Aggregation Server Allocation

In our design, the geo-distributed cloud solution allows us to deploy aggregation servers at many different locations around the world. When a B-box requests to send video data, an aggregation server is assigned to the B-box to aggregate video traffic. Aggregation servers are allocated dynamically according to both the user and original server's locations, the real-time bandwidth load of aggregation servers, and the network performance between the aggregation servers to the user and the original servers. Considering these factors change over time, aggregation servers will be re-allocated regularly, e.g., every 30 minutes or when the network throughput on B-box is low.

We try to maximize the overall bandwidth for all users, while not exceeding the bandwidth capacities of aggregation servers. For example, when two B-boxes request the same aggregation server simultaneously, their original bandwidths of uploading to streaming server using single edge network and the increased bandwidths using aggregated edge networks are often different. Therefore the *bandwidth gain* (i.e., the increased bandwidth comparing to the original bandwidth) are often different too. Suppose this aggregation server's remaining bandwidth capacity can only serve one client, allocating this server to the client with the larger bandwidth improvement can yield more overall bandwidth gain.

Consider all these, we allocate aggregation servers to achieve the following objectives: (1) providing B-boxes with aggregation servers that can achieve high aggregated throughput; (2) matching aggregation servers with B-boxes to achieve maximum overall bandwidth gain.

TABLE I
IMPORTANT NOTATIONS.

Variable	Definition
n	Number of B-boxes requesting AS
m	Number of aggregation servers
t	Current time stamp
C	Set of clients
S	Set of aggregation servers
c_i	B-box i
s_j	Aggregation server j
$B_o(c_i)$	Bandwidth between c_i and original server
$B_o(s_j)$	Bandwidth between s_j and original server
$B(c_i, s_j)$	Bandwidth between c_i and aggregation server s_j
$B_o(c_i, s_j)$	Bandwidth between c_i and original server via s_j
$R(s_j)$	Remaining bandwidths capacity of s_j
$T(s_j)$	Total bandwidths capacity of s_j
$r(s_j)$	Bandwidth load rate of s_j
$G(c_i, s_j)$	Bandwidth gain B-box c_i obtained from s_j
$A(c_i, s_j)$	1 if s_j is allocated to c_i, 0 otherwise
Γ_t	Overall bandwidth gain in time t

B. Algorithm and Implementation

We propose an allocation-matching algorithm to solve the problem above. We summarize important notations in Table I. The allocation-matching algorithm is presented in Algorithm 1.

Suppose a B-box c_i requests for an aggregation server, it retrieves a list of available aggregation servers, uplink bandwidths from these aggregation servers to the original services. Bandwidths between the aggregation servers and the original servers are regularly updated by aggregation servers and report to the scheduling server. Scheduling server preferentially provides B-box with the list of aggregation servers that are close or have better connection to the original server. We define bandwidth load rate $r(s_j)$ as follows:

$$r(s_j) = \frac{R(s_j)}{T(s_j)},$$

Algorithm 1 Aggregation and Matching Algorithm

procedure ALLOCATE SERVER
  for all $c_i \in C$ **do**
    find a subset of S, S_o, which is closest to c
    for all $s_j \in S_o$ **do**
      Measure bandwidth $B(c_i, s_j)$, $B_o(c_i)$
      $B_o(c_i, s_j) \leftarrow min\{B(c_i, s_j), B_o(s_j)\}$
      $G(c_i, s_j) \leftarrow B_o(c_i, s_j) - B_o(c_i)$
      Send bandwidth info to scheduling server
    end for
    Sort $G(c_i)$ in descending order
  end for
  Generate all possible allocation plans satisfy (2)
  Calulate Γ_t for each plan
  return *allocation plan A,*
  that achieves the highest bandwidth gain
  Update $R(s_j)$ for all aggregation server
end procedure

Aggregation servers with $r(s_j)$ lower than a threshold will not be provided to a B-box. The threshold is determined by the bandwidth required to handle user's burst of throughput smoothly. Taking 4K streaming as an example, a stable 15 Mbits/s and max 50 Mbits/s bandwidth is enough to stream 4K. Therefore aggregation servers have to reserve some bandwidth in case of emergency. The threshold can be adjusted according to $T(s_j)$.

B-box then selects a subset S_o of candidate aggregation servers which are in close proximity and measures $B(c_i, s_j)$, which is the uplink bandwidths to these aggregation servers via MPTCP. The bandwidth between c_i and the original server $B_o(c_i)$ normally would be the largest bandwidth among each single edge network. Since without *BASS*, users can only use one edge network. The exact size of this subset and the time spent on measuring uplink bandwidth to each aggregation server is carefully adjusted to avoid adding too much overhead. For example, a B-box can select 3 aggregation servers and spend 5 seconds on measuring the bandwidth to each aggregation server to keep the overall start-up time relatively short.

These bandwidth test results including separate bandwidth on each path are stored in the B-box and updated regularly to avoid measuring bandwidth every time. After the B-box obtains the bandwidth information, it calculates $B_o(c_i, s_j)$, the predicted bandwidth to the original server via these aggregation servers, which is the smaller between the sum of many paths and the bandwidth from aggregation server to original server. Then the B-box reports the bandwidth test results to scheduling server.

Suppose the scheduling server receives n requests from n B-boxes in time t, then it will perform the matching between B-boxes and aggregation servers and try to find a matching that maximizes the overall bandwidth gain Γ_t, as follows:

$$\text{maxmize } \Gamma_t = \sum_{i=1}^{n} \sum_{j=1}^{m} A(c_i, s_j) G(c_i, s_j), \tag{1}$$

subject to

$$\sum_{i=1}^{n} A(c_i, s_j) B_o(c_i, s_j) < R(s_j), j \in \{1, ..., m\} \tag{2}$$

This problem can be solved in a centralized manner. Finally, the scheduling server will allocate aggregation servers according to the calculation and send results back to the B-boxes.

V. EXPERIMENTS AND PERFORMANCE EVALUATION

In this section, we evaluate the performance of our system strategies by real-world experiments on 8 Amazon EC2 instances and 60 Planet-Lab nodes. We begin with the evaluation of the feasibility of selecting nearby aggregation servers to provide high network throughput. Then we evaluate the dynamical allocation of aggregation servers.

A. Experiment Setup

In our experiments, we implemented a B-box prototype on a Raspberry Pi 2 device with kernel version 3.18, which has two Wi-Fi interfaces and runs our own proxy program. We use Amazon EC2 nodes (micro type) to install MPTCP and act as candidate aggregation servers. We use 60 Planet-Lab nodes to act as B-boxes or original servers. The locations of these servers are shown in Fig. 7. Each B-box will be scheduled to upload video data to a randomly selected original server, and an aggregation server will be selected to perform the aggregation and traffic proxy. We deploy bandwidth measurement tools including iperf on the EC2 and Planet-Lab nodes and run the measurements using a Python script.

B. Experiment Results

1) Effectiveness of Distance-based Candidate Aggregation Servers: In Fig. 8, we plot the bandwidth between B-boxes and aggregation servers versus the distance between them. We observe that most PlanetLab servers that obtained higher bandwidth are closer to the Amazon EC2 servers, indicating that the candidate aggregation servers selection step is promising to choose proper aggregation servers.

2) Effectiveness of the Aggregation Server Selection: We compare our algorithm with a random aggregation server selection algorithm, in which aggregation servers are randomly assigned to B-boxes. We define the hit rate γ when selecting an aggregation server as follows:

$$\gamma = \frac{b}{B},$$

where b is the bandwidth achieved using this aggregation server, and B is the optimal bandwidth that can be achieved among all the aggregation servers. $\gamma = 1$ if the algorithm selects the aggregation server with highest bandwidth. The experiment result is shown in Fig. 9. We observe that our

Fig. 7. Locations of testing servers

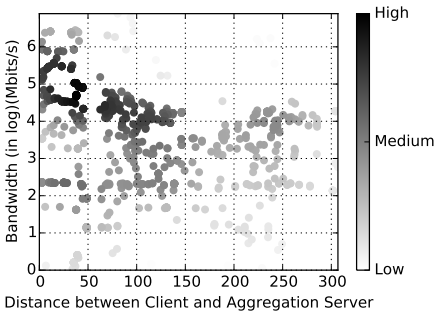

Fig. 8. Relation between bandwidth and distance.

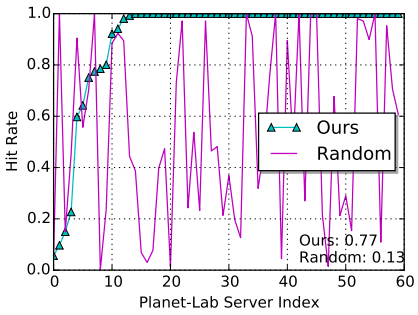

Fig. 9. Comparison of aggregation server allocation.

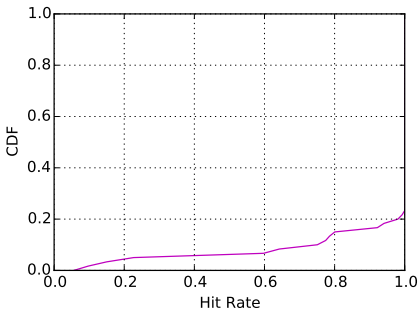

Fig. 10. CDF of Hit Rate γ

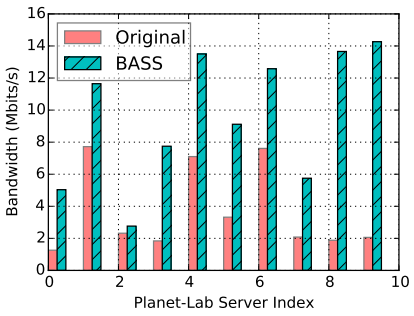

Fig. 11. Bandwidth achieved at a local client

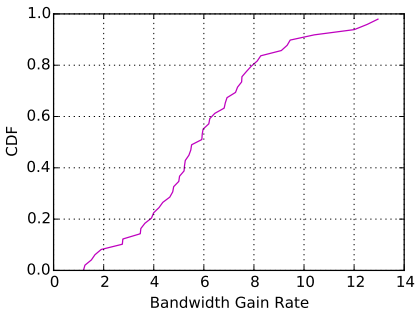

Fig. 12. CDF of bandwidth gain.

algorithm can find the best aggregation servers for the B-boxes and provide the maximum bandwidth most of time. We further plot the CDF of the hit rates in our design in Fig. 10. We observe that for over 77% of the tests, the hit rate γ is 1, indicating that our server allocation performs as well as the optimal solution.

C. Bandwidth Experienced by Clients

Finally, we set up a local PC as the crowdsourced streaming client, which uploads data via B-box prototype. The bandwidths of both interfaces are limited to 8Mits/s. We first measure the bandwidth between client to all possible original servers (i.e., the 60 Planet-Lab nodes). Then we measured the bandwidths achieved by using our BASS system. As illustrated in Fig. 11, the bars are bandwidth achieved in the original TCP paradigm and in BASS, respectively. We observe that BASS system can significantly improve the network performance. The bandwidth can be improved by up to 14 times.

In particular, we measure the improvement in all the tests. We plot the CDF of bandwidth improvement in Fig. 12. The bandwidth gain rate varies between 1.14 and 14.59 with an average of 5.11. For over 50% of the tests, an improvement of 6X is achieved in our design.

VI. Conclusion

With the launch of new crowdsourced live streaming service like Facebook live, its popularity keeps soaring. In this paper, we proposed BASS, a deployable edge network bandwidth aggregation system for improving the uplink bandwidths of broadcasters. We first presented our measurement studies based on a extensive dataset collected by a popular Wi-Fi association mobile app in China. We then presented the BASS architecture, designed an aggregation server allocation strategies, and implemented a prototype on a Raspberry Pi 2 device and Amazon EC2 servers. We evaluated the feasibility and performance on Planet-Lab testbed. The results show stable and promising performance gain. For future work, we plan to implement our system for real-world services, and evaluate its performance in the wild. We plan to (1) extend our system by adding more functions including scheduling video chunk packets according to path characteristics, (2) optimize our system for multiple objectives including cost, Wi-Fi interference and energy consumption.

References

[1] Baochun Li, Zhi Wang, Jiangchuan Liu, and Wenwu Zhu, "Two decades of internet video streaming: A retrospective view," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 9, no. 1s, pp. 33, 2013.

[2] "Periscope hits 10 million users milestone," http://www.slashgear.com/periscope-hits-10-million-users-milestone-13397048/.

[3] Karine Pires and Gwendal Simon, "Dash in twitch: Adaptive bitrate streaming in live game streaming platforms," in *Proceedings of the 2014 Workshop on Design, Quality and Deployment of Adaptive Video Streaming*. ACM, 2014, pp. 13–18.

[4] Cong Zhang and Jiangchuan Liu, "On crowdsourced interactive live streaming: a twitch. tv-based measurement study," in *Proceedings of the 25th ACM Workshop on Network and Operating Systems Support for Digital Audio and Video*. ACM, 2015, pp. 55–60.

[5] Fei Chen, Cong Zhang, Feng Wang, and Jiangchuan Liu, "Crowdsourced live streaming over the cloud," *arXiv preprint arXiv:1502.06314*, 2015.

[6] Ahmed Saeed, Karim Habak, Mahmoud Fouad, and Moustafa Youssef, "Dnis: a middleware for dynamic multiple network interfaces scheduling," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 14, no. 2, pp. 16–18, 2010.

[7] Karim Habak, Moustafa Youssef, and Khaled A Harras, "An optimal deployable bandwidth aggregation system," *Computer Networks*, vol. 57, no. 15, pp. 3067–3080, 2013.

[8] Karim Habak, Moustafa Youssef, Khaled Harras, et al., "G-dbas: a green and deployable bandwidth aggregation system," in *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*. IEEE, 2012, pp. 3290–3295.

[9] Zhen Tang, Zhen Wang, Peng Li, Sini Guo, Xiaofeng Liao, and Hye-Jin Jin, "An application layer protocol for energy-efficient bandwidth aggregation with guaranteed quality-of-experience," 2015.

[10] Alan Ford, Costin Raiciu, Mark Handley, and Olivier Bonaventure, "Tcp extensions for multipath operation with multiple addresses," Tech. Rep., 2013.

[11] Sébastien Barré, Olivier Bonaventure, Costin Raiciu, and Mark Handley, "Experimenting with multipath tcp," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 443–444, 2011.

[12] Costin Raiciu, Christoph Paasch, Sebastien Barre, Alan Ford, Michio Honda, Fabien Duchene, Olivier Bonaventure, and Mark Handley, "How hard can it be? designing and implementing a deployable multipath tcp," in *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association, 2012, pp. 29–29.

[13] Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley, and Hideyuki Tokuda, "Is it still possible to extend tcp?," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 181–194.

[14] Karim Habak, Khaled A Harras, and Moustafa Youssef, "Oscar: A collaborative bandwidth aggregation system," *arXiv preprint arXiv:1401.1258*, 2014.

[15] Gregory Detal, Christoph Paasch, and Olivier Bonaventure, "Multipath in the middle (box)," in *Proceedings of the 2013 workshop on Hot topics in middleboxes and network function virtualization*. ACM, 2013, pp. 1–6.

[16] Yu Wu, Chuan Wu, Bo Li, Linquan Zhang, Zongpeng Li, and Francis C.M. Lau, "Scaling Social Media Applications into Geo-Distributed Clouds," in *IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2012.

[17] Y Angela Wang, Cheng Huang, Jin Li, and Keith W Ross, "Estimating the performance of hypothetical cloud service deployments: A measurement-based approach," in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 2372–2380.